



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/721,316	11/25/2003	Daniel Citron	IL920030047US1	7050

7590 09/13/2006

Stephen C. Kaufman
IBM Corporation
Intellectual Property Law Dept.
P.O. Box 218
Yorktown Heights, NY 10598

EXAMINER

TSAI, SHENG JEN

ART UNIT	PAPER NUMBER
----------	--------------

2186

DATE MAILED: 09/13/2006

Please find below and/or attached an Office communication concerning this application or proceeding.

Office Action Summary

Application No.

10/721,316

Applicant(s)

CITRON, DANIEL

Examiner

Sheng-Jen Tsai

Art Unit

2186

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) OR THIRTY (30) DAYS, WHICHEVER IS LONGER, FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

- 1) ☒ Responsive to communication(s) filed on 22 August 2006.
- 2a) ☒ This action is **FINAL**. 2b) ☐ This action is non-final.
- 3) ☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

- 4) ☒ Claim(s) 1-26 is/are pending in the application.
- 4a) Of the above claim(s) 13 and 26 is/are withdrawn from consideration.
- 5) ☐ Claim(s) _____ is/are allowed.
- 6) ☒ Claim(s) 1-12 and 14-25 is/are rejected.
- 7) ☐ Claim(s) _____ is/are objected to.
- 8) ☐ Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

- 9) ☐ The specification is objected to by the Examiner.
- 10) ☐ The drawing(s) filed on _____ is/are: a) ☐ accepted or b) ☐ objected to by the Examiner.
Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).
Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).
- 11) ☐ The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

- 12) ☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).
- a) ☐ All b) ☐ Some * c) ☐ None of:
- ☐ Certified copies of the priority documents have been received.
 - ☐ Certified copies of the priority documents have been received in Application No. _____.
 - ☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

* See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

- | | |
|-----------------------------------------------------------------------------------------|-----------------------------------------------------------------------------|
| 1) <input type="checkbox"/> Notice of References Cited (PTO-892) | 4) <input type="checkbox"/> Interview Summary (PTO-413) |
| 2) <input type="checkbox"/> Notice of Draftsperson's Patent Drawing Review (PTO-948) | Paper No(s)/Mail Date. _____ |
| 3) <input type="checkbox"/> Information Disclosure Statement(s) (PTO-1449 or PTO/SB/08) | 5) <input type="checkbox"/> Notice of Informal Patent Application (PTO-152) |
| Paper No(s)/Mail Date _____ | 6) <input type="checkbox"/> Other: _____ |

DETAILED ACTION

1. This Office Action is taken in response to Applicant's Amendments and Remarks filed on August 22, 2006 regarding application 10,721,316 filed on November 25, 2003.

2. Claims 13 and 26 have been cancelled.

Claims 1-12, 14-17 and 24 have been amended.

Claims 1-12 and 14-25 are pending for consideration.

3. ***Response to Amendments and Remarks***

Applicant's amendments and remarks have been fully and carefully considered, with the Examiner's response set forth below.

integrated circuit chip vs. integrated circuit device

Applicants amends the claims to recite that the invention being in the form of "an integrated circuit chip" instead of being in the form of "an integrated circuit device," and contends that the prior art (Citron et al., "Creating a Wider Bus Using caching Techniques," Proceedings of the First IEEE Symposium on High-Performance Computer Architecture, January, 1995, pages 90-99) does not teach the limitation that all components (processing element, L1 cache, L2 cache, buses and bus expanders) reside on the same chip (i.e., on-chip). The examiner disagrees with this assessment for the following reasons:

First, Citron et al. teach that the processing element and the bus expanders are on the same chip [when used to reduce the number of I/O pins on a microprocessor (i.e., the corresponding processing element), the bus expander should reside on the chip itself (page 91, Section 2, Description, first paragraph)].

Second, Citron et al. teach that the cache be on-chip [page 95, Section 3.2.2, On-chip cache present; page 95, Section 3.2.4, On-chip cache write policy].

Third, Citron et al. teach the use of on-chip L1 and L2 caches [page 96, Section 3.2.5, On-chip cache parameters; page 97, "Table 5 indicates that the same level of compaction is attained both between a L1 cache to memory (or L2 cache) and between a L2 cache to memory" (page 97, left column, first paragraph)].

Fourth, since the processing element, L1 and L2 caches and bus expanders are all on-chip, it follows that all the buses connecting between these elements are also on-chip.

Therefore, Citron et al. do teach that all components being on-chip and form an integrated circuit chip.

Another iteration of claim analysis based on the previously identified reference (Citron et al., "Creating a Wider Bus Using caching Techniques," Proceedings of the First IEEE Symposium on High-Performance Computer Architecture, January, 1995, pages 90-99) and addressing newly amended as well as previously presented limitations of the claims has been embarked. Refer to the corresponding sections of the claim analysis for details.

Claim Rejections - 35 USC § 102

4. The following is a quotation of the appropriate paragraphs of 35 U.S.C. 102 that form the basis for the rejections under this section made in this Office action:

A person shall be entitled to a patent unless –

(b) the invention was patented or described in a printed publication in this or a foreign country or in public use or on sale in this country, more than one year prior to the date of application for patent in the United States.

5. Claims 1-26 are rejected under 35 U.S.C. 102(b) as being anticipated by Citron et al. ("Creating a Wider Bus Using caching Techniques," Proceedings of the First IEEE Symposium on High-Performance Computer Architecture, January, 1995, pages 90-99).

As to claim 1, Citron et al. disclose **an integrated circuit chip** [Citron et al. teach that all components being on-chip and form an integrated circuit chip (refer to "integrated circuit chip vs. integrated circuit device" presented earlier in this Office Action)], **comprising:**

a processing component on the integrated circuit chip [figures 1 and 2 show the components on the chip; the bus-expander is a component that mediates between a device with digital input and output such as a processor (page 91, section 2, Description, first paragraph); when used to reduce the number of I/O pins on a microprocessor (i.e., the corresponding processing element), the bus expander should reside on the chip itself (page 91, Section 2, Description, first paragraph)],

a Level 1 (L1) cache and a Level 2 (L2) cache on the chip [page 96, Section 3.2.5, On-chip cache parameters; page 97, "Table 5 indicates that the same level of compaction is attained both between a L1 cache to memory (or L2 cache) and between a L2 cache to memory" (page 97, left column, first paragraph)], **which are arranged to store data for use by the processing component responsively to an addressing scheme** [this is the inherent property of a cache memory] **based on memory addresses having an address length of m_1 bits** [the bus expander maps an m-bit wide value into a smaller, n-bit wide value (page 91, section 2, Description, first paragraph); the information/value may be instructions, addresses or data (page 91,

section 2, Description, second paragraph); compact the high-order bits of an address into a small number of bits representing the index of a cache-like device of the high-order bits of recently accessed addresses (page 90, right column, second paragraph)); **and**

first and second buses [the bus expander having a split (Harvard) architecture so that data, addresses, instructions and instruction addresses are all stored separately (page 94, first paragraph), this necessitates the first (for data) and second (for address) buses; there are many needs for compaction of both data and address values (page 91, first paragraph); most devices and buses have the same bit width for address and data, ..., it is possible to store address and data values in the same LUT (page 92, second paragraph); thus the first and second buses refer to the address and data buses, respectively, connecting the processing element and the L1 cache **on the chip** [since the processing element, L1 and L2 caches and bus expanders are all on-chip, it follows that all the buses connecting between these elements are also on-chip (refer to "*integrated circuit chip vs. integrated circuit device*" presented earlier in this Office Action)] **coupled between the processing component and the L1 cache** [compaction/expansion units can be placed between processor and memory (abstract); it is well known in the art that L1 cache is connected directly to the processing unit], **and third and fourth buses on the chip coupled between the L1 cache and the L2 cache** [compaction/expansion units can be placed devices that access system bus (abstract); it is well known in the art that the L2 cache is connected directly to L1 cache and provides data upon a "miss" of L1 cache; the third and fourth buses are the

address and data buses, respectively, connecting between the L1 and L2 caches, and are the counterparts of the first and second buses], **the first and third buses having bus widths smaller than m_1** [the bus expander maps an m -bit wide value into a smaller, n -bit wide value (page 91, section 2, Description, first paragraph), thus $n < m$], **the processing component and the L1 and L2 caches comprising respective address bus expanders coupled to the first and the third buses** [figures 1 and 2 illustrate this configuration; the first and second buses refer to the address and data buses connecting the processing element and the L1 cache; the third and fourth buses are the address and data buses connecting between the L1 and L2 caches; compaction/expansion units can be placed between processor and memory, between processor and local bus, and between devices that access the system bus (abstract)] **in order to compact at least at least some memory addresses is transmitted over the first and the third buses** [compact the high-order bits of an address into a small number of bits representing the index of a cache-like device of the high-order bits of recently accessed addresses (page 90, right column, second paragraph)] so that each **of the at least some memory addresses is transmission over the first and the third buses in one bus cycle** [a bus expander translates a device-word that is to be communicated over a bus into information that can be transferred over the bus in a single bus cycle (page 91, section 2, Description, second paragraph)].

Wherein the L1 cache comprises a table [the main component of the bus expander is a look up table (LUT) (page 91, right column, first paragraph)] **for use in compaction of address fields that is shared by the address bus expanders**

coupled to the first and third buses [figure 10 gives a breakdown of the hit ratio on the address-LUT and data-LUT (page 96, left column, first paragraph), which means that both the address and data are handled simultaneously; a cycle hit occurs if, during a write cycle, both the address-LUT and data-LUT registered hits and a cycle miss occurs if either one registered a miss (page 93, right column, first paragraph); most devices and buses have the same bit width for address and data, ..., it is possible to store address and data values in the same LUT (page 92, second paragraph)].

As to claim 2, Citron et al. teach that **the data comprise data words having a word length of m_2 bits stored at each address, such that the second bus has a bus width smaller than m_2 , and wherein each of the processing component and the cache further comprises a respective second bus expander coupled to the second bus in order to compact at least some of the data words for transmission over the second bus so that each of the at least some of the data words is transmitted over the second bus in one cycle of the second bus** [the data portion is handled in exactly the same way as the address is handled, as explained in "As to claim 1;" our measurements show that locality is also present with data value (page 90, right column, third paragraph); the bus expander having a split (Harvard) architecture so that data, addresses, instructions and instruction addresses are all stored separately (page 94, first paragraph), this necessitates the first (for data) and second (for address) buses; the bus expander maps an m-bit wide value into a smaller, n-bit wide value (page 91, section 2, Description, first paragraph), thus $n < m$; a bus expander translates a device-word that is to be communicated over a bus into

information that can be transferred over the bus in a single bus cycle (page 91, section 2, Description, second paragraph)].

As to claim 3, Citron et al. teach that **the data words comprise data values for processing by the device, and wherein the processing component is arranged to load the compacted data words via the second bus from the cache for processing and to store the compacted data words via the second bus to the cache** [the data portion is handled in exactly the same way as the address is handled, as explained in "As to claim 1;" our measurements show that locality is also present with data value (page 90, right column, third paragraph); the bus expander having a split (Harvard) architecture so that data, addresses, instructions and instruction addresses are all stored separately (page 94, first paragraph), this necessitates the first (for data) and second (for address) buses; the bus expander maps an m-bit wide value into a smaller, n-bit wide value (page 91, section 2, Description, first paragraph), thus $n < m$; a bus expander translates a device-word that is to be communicated over a bus into information that can be transferred over the bus in a single bus cycle (page 91, section 2, Description, second paragraph); section 3.2.2 discusses "on-chip cache present;" compact the high-order bits of an address into a small number of bits representing the index of a cache-like device of the high-order bits of recently accessed addresses (page 90, right column, second paragraph)].

As to claim 4, Citron et al. teach that **the data words comprise instructions for execution by the device, wherein the compacted words comprise compacted instructions, and wherein the processing component is arranged to fetch the**

compacted instructions via the second bus [the instruction portion is handled in exactly the same way as the address/data is handled, as explained in "As to claim 1;" the bus expander having a split (Harvard) architecture so that data, addresses, instructions and instruction addresses are all stored separately (page 94, first paragraph), this necessitates the first (for data) and second (for address) buses; the bus expander maps an m-bit wide value into a smaller, n-bit wide value (page 91, section 2, Description, first paragraph), thus $n < m$; a bus expander translates a device-word that is to be communicated over a bus into information that can be transferred over the bus in a single bus cycle (page 91, section 2, Description, second paragraph); section 3.2.2 discusses "on-chip cache present;" compact the high-order bits of an address into a small number of bits representing the index of a cache-like device of the high-order bits of recently accessed addresses (page 90, right column, second paragraph)].

As to claim 5, Citron et al. teach that **the address bus expander and the second bus expander are arranged to compact the memory addresses and the data words simultaneously, so as to transmit a compacted memory address and a compacted data word for storage at the memory address together in one cycle of the first and second buses** [figure 10 gives a breakdown of the hit ratio on the address-LUT and data-LUT (page 96, left column, first paragraph), which means that both the address and data are handled simultaneously; a cycle hit occurs if, during a write cycle, both the address-LUT and data-LUT registered hits and a cycle miss occurs if either one registered a miss (page 93, right column, first paragraph)].

As to claim 6, Citron et al. teach that **the address bus expander and the second bus expander are arranged to compact the memory addresses and the data words by transmitting indices to values in respective tables held by the bus expanders, and wherein the cache is arranged to store at least some of the indices together with the data** [compact the high-order bits of an address into a small number of bits representing the index of a cache-like device of the high-order bits of recently accessed addresses (page 90, right column, second paragraph); during compaction, part of the device-word is the input of the LUT and the output is an index of the LUT where device-word is contained (page 91, right column, first paragraph)].

As to claim 7, Citron et al. teach that **the address bus expander is arranged to compact each of the at least some of the memory addresses by dividing each of the memory addresses into at least first and second fields** [during compaction, the LUT is search for the high order bits by dividing these bits into two fields, a tag and a key (page 92, left column, first paragraph)], **storing values of the second field in a respective table such that the values in respective tables held by the address bus expander in the processing component and the address bus expander in the cache are identical, and if the second field of a memory address matches a value in the table, transmitting an index corresponding to the value over the first bus along with the first field in the one cycle of the bus** [the components of the set in this LUT line are searched for a matching tag. If successful, the output is the key and the set number where the tag is found (page 92, left column, first paragraph)].

As to claim 8, Citron et al. teach that **the first field comprises a set of least significant bits (LSB) of the memory address, while the second field comprises a set of most significant bits (MSB) of the memory address** [compact the high-order bits of an address into a small number of bits representing the index of a cache-like device of the high-order bits of recently accessed addresses (page 90, right column, second paragraph); the key, tag and low order bits are assembled to form the device-word (page 92, left column, first paragraph)].

As to claim 9, Citron et al. teach that **the at least first and second fields comprise a third field, and wherein the address bus expander is arranged to compact each of the at least some of the memory addresses by transmitting first and second indices corresponding to the values of the first and third fields, respectively, over the first bus along with the first field** [the key, tag and low order bits are assembled to form the device-word (page 92, left column, first paragraph)].

Note that the key and the tag belong to the high order bits field].

As to claim 10, Citron et al. teach that **the address bus expander in the processing component is arranged, when the second field of the memory address does not match any of the values in the table, to transmit both of the first and second fields over multiple cycles of the bus, and to cause the respective table of the bus expander to be updated in both the processing component and the cache** [when a compaction fails, the high order bit indicates this fact. The LUT is updated so future references will not fail and the entire device-word

must be passed to the receiver (page 92, left column, section 2.1, Processing a Miss, first paragraph)].

As to claim 11, Citron et al. teach that **the cache comprises lines of the data that are indexed according to the first field, each line containing a corresponding value of the second field, and wherein the address bus expander in the cache is arranged, upon receiving the index over the first bus [compact the high-order bits of an address into a small number of bits representing the index of a cache-like device of the high-order bits of recently accessed addresses (page 90, right column, second paragraph)], to retrieve the value of the second field from the table responsively to the index [the components of the set in this LUT line are searched for a matching tag. If successful, the output is the key and the set number where the tag is found (page 92, left column, first paragraph)], and wherein the cache is arranged to determine whether a cache hit has occurred by checking the retrieved value against the corresponding value of the second field in the line that is indexed by the first field [a cycle hit occurs if, during a write cycle, both the address-LUT and data-LUT registered hits and a cycle miss occurs if either one registered a miss (page 93, right column, first paragraph)].**

As to claim 12, Citron et al. teach that **the address bus expander is arranged to retrieve the value of the second field from the table simultaneously with retrieval of the data from the line in the cache that is indexed by the first field for transmission of the data over the second bus to the processing component [a cycle hit occurs if, during a write cycle, both the address-LUT and data-LUT registered**

hits and a cycle miss occurs if either one registered a miss (page 93, right column, first paragraph); figure 10 gives a breakdown of the hit ratio on the address-LUT and data-LUT (page 96, left column, first paragraph), which means that both the address and data are handled simultaneously; there are many needs for compaction of both data and address values (page 91, left column, first paragraph)).

As to claim 14, refer to "As to claim 1" presented earlier in this Office Action.

As to claim 15, refer to "As to claim 2" presented earlier in this Office Action.

As to claim 16, refer to "As to claim 3" presented earlier in this Office Action.

As to claim 17, refer to "As to claim 4" presented earlier in this Office Action.

As to claim 18, refer to "As to claim 5" presented earlier in this Office Action.

As to claim 19, refer to "As to claim 6" presented earlier in this Office Action.

As to claim 20, refer to "As to claim 7" presented earlier in this Office Action.

As to claim 21, refer to "As to claim 8" presented earlier in this Office Action.

As to claim 22, refer to "As to claim 9" presented earlier in this Office Action.

As to claim 23, refer to "As to claim 10" presented earlier in this Office Action.

As to claim 24, refer to "As to claim 11" presented earlier in this Office Action.

As to claim 25, refer to "As to claim 12" presented earlier in this Office Action.

Conclusion

6. Claims 1-12 and 14-25 are rejected as explained above.

7. **THIS ACTION IS MADE FINAL.** Applicant is reminded of the extension of time policy as set forth in 37 CFR 1.136(a).

A shortened statutory period for reply to this final action is set to expire THREE MONTHS from the mailing date of this action. In the event a first reply is filed within

TWO MONTHS of the mailing date of this final action and the advisory action is not mailed until after the end of the THREE-MONTH shortened statutory period, then the shortened statutory period will expire on the date the advisory action is mailed, and any extension fee pursuant to 37 CFR 1.136(a) will be calculated from the mailing date of the advisory action. In no event, however, will the statutory period for reply expire later than SIX MONTHS from the mailing date of this final action.

8. Any inquiry concerning this communication or earlier communications from the examiner should be directed to Sheng-Jen Tsai whose telephone number is 571-272-4244. The examiner can normally be reached on 8:30 - 5:00.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Matthew Kim can be reached on 571-272-4182. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see <http://pair-direct.uspto.gov>. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free).

Sheng-Jen Tsai
Examiner
Art Unit 2186

August 31, 2006



MATTHEW KIM
SUPERVISORY PATENT EXAMINER
TECHNOLOGY CENTER 2100